AD-A125 696 SPATIAL AND TEMPORAL VISUAL MASKING AND VISIBILITY(U) 1/1
NEW HAMPSHIRE UNIV DURHAM VISION RESEARCH LAB
R A SMITH 01 OCT 82 AFOSR-TR-83-0050 AFOSR-80-0045

UNCLASSIFIED F/G 5/10 NL

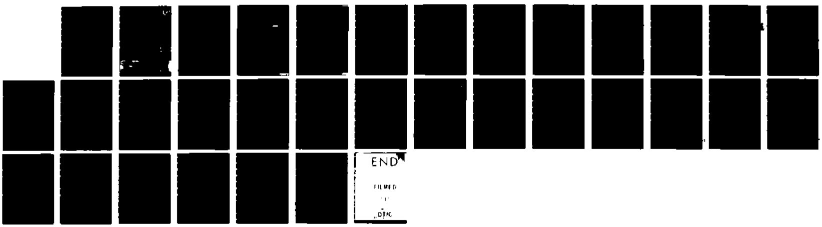

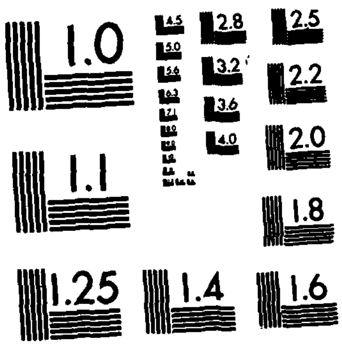

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AFOSR-TR- 83-0050

VISION RESEARCH LAB

SPATIAL AND TEMPORAL VISUAL MASKING AND VISIBILITY

First Annual Report
1 October 1982

University [illegible] shire

SPATIAL AND TEMPORAL VISUAL MASKING AND VISIBILITY

First Annual Report
1 October 1982

For: Air Force Office of Scientific Research

Prepared by

ROBERT A. SMITH

Vision Research Laboratory
of UNH
Pettee Brook Offices
Durham, NH 03824

AFOSR-80-0045

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER **AFOSR-TR- 83-0050**	2. GOVT ACCESSION NO. AD-A125 694	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Spatial And Temporal Visual Masking And Visibility		5. TYPE OF REPORT & PERIOD COVERED Annual Report FY '82
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Robert A. Smith, PhD		8. CONTRACT OR GRANT NUMBER(s) AFOSR-80-0045
9. PERFORMING ORGANIZATION NAME AND ADDRESS Vision Research Lab of UNH, Pette Brook Offices, Durham, NH 03824		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 61102F [redacted] 2313/A5
11. CONTROLLING OFFICE NAME AND ADDRESS AFOSR/NL, Building 410, Bolling AFB, DC 20332		12. REPORT DATE ~~8/1/82~~
		13. NUMBER OF PAGES 29
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Spatial, Masking, Visibility, Vision, Spatial Frequency

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

In the past year work has progressed in four interrelated areas.
1) Substantial evidence supports the hypothesis that visual detection in the presence of masking noise occurs at a constant signal/noise ratio only if the subject is unfamiliar with the mask.
2) A variety of studies of the proposed sustained/transient mechanism model have failed to find support for a dichotomous visual system, and have failed to replicate some of the classical studies in this area. (over)

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

Item 20 continued:

Studies of hypothetical velocity-tuned channels suggest the possible existence of such channels, but show that they are, at best, mixed in their tuning properties.

Weisstein'g Pixel Flicker technique has been found to be an artifact of an unrealistic algorithm for adding noise to a visual scene. Other temporal enhancement schemes show some promise.

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

b. Objectives

Overall -- Study the interraction of spatial and temporal variations on stimulus detectability.

1. Further elucidate our hypothesis that different detection criteria amy be used in masking discriminations, and that these produce qualitatively different threshold behavior.

2. Study the detectability of spatially and temporally varying sinusoids, with the ultimate aim of testing the familiar sustained and transient mechanisms hypothesis.

3. Study various quantitative measures of velocity tuning in visual channels.

4. Set up our image processor and evaluate some temporal algorithms for image enhancement.

c. Status of the Research

1. Introduction

This has been the first year in which members of this project have worked essentially full time. Despite some significant time losses (e.g. the principle investigators absence because of illness) the results have been most encouraging. We currently have seven papers in various stages of publication and a variety of pilot studies and new projects underway.

A considerable portion of the current year has been devoted to practical considerations. At the beginning of the Grant year our lab was moved from the engineering building to a suite of offices adjacent to the UNH campus. Despite the time involved, the move has proven advantageous in that our new space is substantially more efficient than the old. A major problem arose in April when Digital Equipment announced it's intention of reposssessing the PDP 11 computer which runs our Grinnel Visual Display. Since work on the visual display was at the software development stage, this uncertainty about what computer we would ultimately be using effectively halted all further work on the display. After six months of intense negotiation, Digital Equipment finally decided to donate the computer to our project. We are pleased to have the computer, but the loss of six months work is unfortunate.

2. Status at the Beginning of the Reported Year

The conclusion of our previous grant left us with some, but not all, of the directions of this work firmly established. Our hypothesized mechanisms for masking detection were worked out but much elaboration and data collection remained. The data on sustained and transient detection (Steve Panish's thesis) were partly collected, and the general outlines of these results were becoming clear. Several years ago Dan Swift had begun working on the idea that if there are separate systems for detection of form and motion, then it would be appropriate to abandon pure detection tasks in favor of tasks which are clearly specific to one particular system. For example, spatial frequency discrimination would seem to be specific to the form system. He then took some preliminary data on the effect of stimulus motion on such tasks, shedding some tentative light on motion tuning of purely spatial channels. Finally we had some pilot data on the enhancement of images by subthreshold enhancers and by pixel flicker. All of these were done in one dimension, however.

2. Status of our Research in the Current Year

aa. Brief Summary

We have substantially strengthened our hypothesis that Weber's Law in spatial masking occurs when the subject is unable to predict the appearance of the mask pattern. In addition to memorizing the mask pattern, we find that subjects can predict the appearance of the mask pattern under three-alternative forced-choice conditions, even if the mask is changed on every trial. This is a strong prediction of our hypothesis. Recently we have turned our attention to other psychophysical methods, particularly method of adjustment. We find that a generalization of Birdsall's Theorem (Lasley and Cohn, 1981) provides a plausible explanation for the common observation of Weber's Law with method-of-adjustment. We hypothesize that subjects use, as a criterion for detection, some specific change in the stimulus pattern (e.g. threshold might be a 10% change in perceived mask bar width). It is readily shown that such a criterion occurs when mask and test bear a fixed ratio to each other; thus test threshold will be proportional to mask modulation, which implies Weber's Law. We find that it is possible for a subject to use different criterion in setting adjustment thresholds. If he uses a fixed configurational criterion (as above) then Weber's Law is observed. However, a more sensitive criterion (e.g. ANY discernable change in the mask pattern) yields power law masking.

Our study of spatial and temporal mechanisms has continued after the completion of Steve Panish's Phd. dissertation. The findings of this dissertation include 1. flicker and pattern threshold functions were studied using both counterphase and moving gratings. The form of these functions is different with the two types of gratings, which we tentatively attribute to involuntary eye movements. We find that differences between flicker and pattern thresholds are very much smaller than has been reported in the past. 2. The time course and recovery of spatial adaptation were measured under conditions designed to isolate the sustained and transient systems. No characteristic time-course could be found for either system. 3. The effect of spatio-temporal adaptation on the perception of pattern and flicker was measured with a bias-correcting rating procedure using near threshold counterphase flicker gratings. False alarm rates were much more important in the ratings of flicker than pattern. When false alarms were subtracted from the rating scores, we found that the sensation of pattern predominated at threshold. We hypothesize that the reported findings of lower thresholds for flicker than pattern may represent a problem with criterion shift.

Several experiments have been used to assess the velocity tuning of spatial channels. In one experiment we measured the magnitude of the spatial-frequency shift and variations in apparent contrast of the test field as a function of adapt and test velocity. We obtained very clean data showing that the magnitude of the spatial-frequency shift is not tuned for velocity, but is generally enhanced by nonzero adaptation velocities. The enchancement occurs irrespective of direction of motion of the adapt and test gratings. This result suggests that velocity plays only a very simple role in a task that is purely spatial. Unfortunately the contrast-matching data (which we hope will parallel these results) are still to noisy for certainty.

In another substantial experiment on motion tuning, we studied the magnitude of the motion after-effect as a function of adapt and test motion. The results of this are complex, and no simple generalization describes them. It is quite clear, however, that the MAE is not at all a simple additive velocity, as has sometimes been claimed. The hypothesis which most closely accounts for our data is that the MAE is a velocity-shift phenomenon, analogous to the spatial frequency shift. Superimposed upon this, however, there appears to be increased AE with increased adapt velocity. There is some indication of velocity tuning as well, but other effects are probably superimposed upon it. The motion system may be much more complex than the spatial one.

In the past six months, we have set up the image processor to study the phenomenon of Pixel Flicker. We replicated the well-known demonstration by Genter and Weisstein (OSA meeting, Sarasota, 1980), and found that it indeed produced a striking improvement in image quality. The most impressive aspect of this demonstration is the fact that considerable image detail emerges from a noisy image, which initially shows little or no detail. Unfortunately, we quickly determined that this is largely artifactual. Genter and Weisstein used an unrealistic alogrithm for adding noise to their image; noise with a uniform distribution was added to each pixel using modulo 256 arithmetic. (In a realistic situation, overflows would be clamped at 255.) It can be shown that the use of modulo-noise degrades the image much more than clamped noise. Thus our demonstation that pixel flicker improves detection at low spatial frequencies is still valid. However the likelihood that pixel flicker is qualitatively more effective with two-dimensional images, as Genter and Weisstein claimed, now seems very small. We are exploring this and other dynamic image modifications to see which appear qualitatively promising, though it will be some time before we are able to do quantitative measurements of detectability.

In the following sections, we consider all of these projects in detail.

bb. Weber's Law, Masking, and Birdsall's Theorem

For some years there has been a controversy in the field of spatial frequency masking over whether masking obeys Weber's Law or whether it obeys a power law. That is, does the threshold for the test grating rise proportionally to the contrast of the mask grating, or does it rise proportional to some power (less than 1) of the mask contrast. Two years ago (Smith and Swift, 1980), we discussed -- and ultimately disproved -- three hypotheses about Weber's Law which have been proposed in the literature. Now we present our own hypothesis which we believe explains all of the evidence currently available to us. We propose that the subject in a masking experiment has different strategies available to him for setting thresholds, and that these different strategies yield qualitatively different behavior (i.e. Weber's Law or power law). Different experimental paradigms will favor one strategy over another, accounting for the fact that some experimenters have found one or another of the laws to hold, and that a few groups (such as our own) have reported both.

The essence of one of our detection strategies is described in a theorem proposed by Lasley and Cohn (1981), and originally attributed to Birdsall. Although Lasley and Cohn applied Birdsall's Theorem to an identification task, we find it equally applicable to detection in the presence of spatial masking. Birdsall's Theorem states that if the limiting constraint on the detection of a visual signal is noise which is external to the visual system, and if the observer uses an optimal strategy, than any monotonic, nonlinear transformation which the visual system may impose upon its input will have no detectable effect upon threshold, and threshold will be proportional to the contrast of the masking noise (i.e. Weber's Law). This suggests to us that those studies which find Weber's Law to hold may be operating under conditions such that Birdsall's theorem applies. Conversely we suggest that those studies which fail to find Weber's Law are operating under such conditions that Birdsall's Theorem is inapplicable, and that the resulting nonlinear behavior represents real nonlinearity in the visual system. In a sense, those experiments which find a power law are more informative to students of visual physiology.

Figure 1 shows a conceptual model of the visual system, in terms of which we will view our results. In this figure we see the outputs of a variety of spatial channels of different frequency while viewing a pattern of random noise, which may or may not have a test grating added to it. Also indicated are the mean output of all the channels (I_{mean}), a measure (σ) of the variation of these outputs, and I, the output of the channel most sensitive to the test grating. The question of whether or not a signal is detectable in this pattern of channel responses reduces to the question of whether I is sufficiently large that it is unlikely to have occurred by chance in the random mask. This is a simple question of statistics. Allowing the visual system some knowledge of the

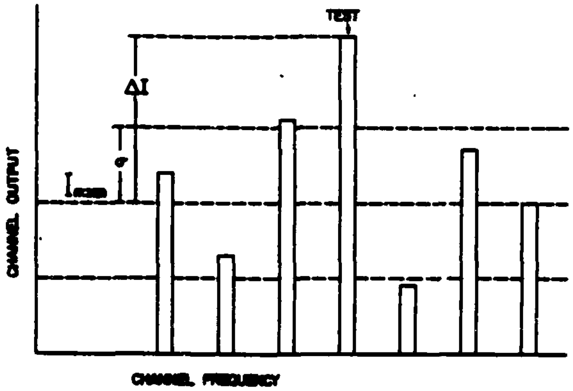

Fig. 1. A conceptualization of our model of the visual system. This shows the outputs of 7 channels of varying center frequencies, with the outputs varying around a mean value, Imean. The channel most sensitive to the test frequency has an output, relative to Imean, of ΔI. σ is a measure of the variance of the outputs of the channels.

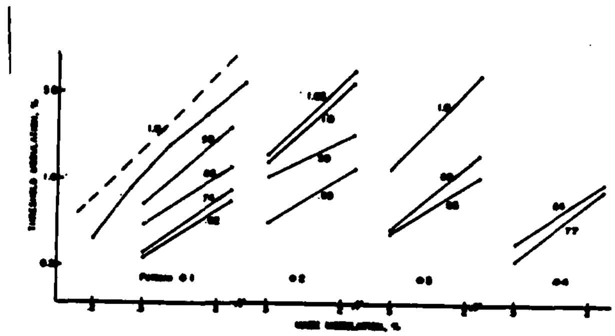

Fig. 2. Test threshold versus mask contrast, for 4 different noise mask patterns. Successive trials with the same mask are plotted below their predecessors. The decrease in threshold is real; these data have not been displaced for clarity.

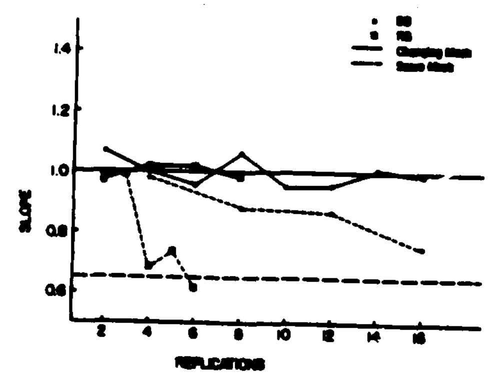

Fig. 3. The slope of the masking function (e.g. fig 2) versus number of trials practice. Solid lines are for the case where the mask was changed on every trial, dotted lines are for when the mask was not changed

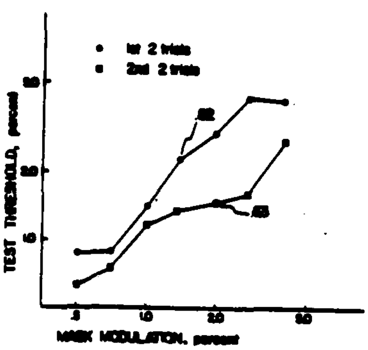

Fig. 4. Similar to Fig. 3, except that 3-alternative forced choice psychophysics were used, with changing masks.

distribution of channel outputs, I is simply compared to the width of this distribution (σ) by means of a critical ratio (I/σ). This is entirely analogous to the familiar T-test in statistics. We note that is proportional to the intensity of the pattern (I_{mean}) so that the critical ratio is proportional to $\Delta I/I_{mean}$, from which we see that Weber's Law applies to masking in the presence of external noise. The fact that nonlinear transformations within the visual system have no effect follows from the observation that the constraints on detection lie within the stimulus itself, and are already effective before the stimulus even enters the visual system. A monotonic nonlinearty in the visual system may change the shape of the distribution of channel outputs, but so long as the observer uses an optimal strategy (which means, in effect, that he has adequate knowledge of the shape of this distribution) then the threshold dictated by the critical ratio remains unchanged. Note that the condition that the subject use an optimal detection strategy implies that under some conditions a period of learning may be necessary. We have generally found this to be the case.

Let us now consider some data showing a subject learning to distinguish the presence of a four cycle/degree test grating in the presence of broad-band one dimensional visual noise. (Our noise stimuli were broad-band white noise, with a logarithmic amplitude density over the range of 2-8 cycles/degree. In practice, these were generated by adding together 8 sinsusoids of equal amplitude and randomly chosen phase and with frequencies spaced at equal logarithmic intervals across the of 2-8 cycles/degree. By changing the phases of the constituant gratings, we could generate a variety of noise patterns with identical power spectra -- apart from edge effects -- but with very different appearance.) The upper left hand line in Figure 2 shows the subject's thresholds when first exposed to this particular noise mask. It will be seen that the curve relating threshold to mask contrast is linear and has a slope of 1.0; i.e. Weber's Law applies here. We now select two convenient points along this curve, and repeatedly measure the subject's threshold at these points. It will be seen that the threshold decreases monotonically with experience (these curves have not been displaced for clarity) but more important we find that the slope of the masking function decreases with practice to an asymptote of about 0.65. The second column of Figure 2 shows what happens when the subject is presented with a new mask pattern. Initially the slope of the masking function rises to 1, but again, with practice, threshold decreases and the slope drops to about .65. The experiment is repeated in the remaining columns of Figure 2. From the results of this experiment, we may conclude that the very same masking stimulus may yield Weber's Law or a power law under different conditions; in this experiment, the relevant condition seems to be familiarity with the mask stimulus.

If familiarity with the noise mask is indeed the relevant condition, then it should be possible to prevent the subject from developing a power law by changing the mask stimulus on each and every trial. The results of such an experiment are shown in Figure 3, which plots the slope as a function of trials, rather than the raw data. Data from the previous experiment are also included. Although both experiments start with a slope of 1, the data for the continually changing noise pattern retain this slope indefinitly while data for the same noise pattern used repeatedly quickly fall to a slope of about .65.

The experiments presented thus far may be interpreted in the following way. When a mask is first presented to a subject, it is unfamiliar and is treated as noise in the sense of Birdsall's Theorem. With repeated exposure, however, the observer memorizes the appearance of the noise mask, it ceases to be an unknown pattern, and so ceases to function as noise. This is in good agreement with the subjective report of the subjects, who set thresholds in a different way in the two situations. If the noise mask is unfamiliar the subject must actually look for the presence of a periodicity at the frequency of the test stimulus. When the mask is familiar, however, the subject looks for small deviations from an otherwise well-remembered pattern.

There is another way in which we can provide the subject with sufficient information to distinguish test-plus-mask from mask-along. In the previous experiments, we believe that the well-rehearsed subject was comparing the two presented gratings with a third grating with was the memorized image of the mask. We can provide all of this information in a single trial by switching from two-alternative to three-alternative forced choice. In this paradigm, the subject is presented with two mask gratings and one mask-plus-test grating; his task is to pick the odd grating. If our hypothesis is correct, it is clear that even with an unfamiliar mask the subject should be able to discriminate in a power law fashion. Figure 4 shows that this is the case. Here the subject is doing three-alternative forced choice, with a different mask used on each and every trial. Nonetheless, these data clearly display a power law, something which never occurs with two-alternative forced choice and changing masks. These data are particularly significant because they were taken from an entirely naive and unpracticed subject. Masking discriminations are among the most difficult in visual psychophysics, and to the best of our knowledge, these are the first data reported on subjects who have not had extensive practice.

The data thus far presented are all quite consistant with the hypothesis that Birdsall's Theorem applies to spatial frequency masking when the mask is an unfamiliar visual noise pattern. However the majority of spatial frequency masking experiments are done not with a noise mask but with an harmonically pure mask. We offer two suggestions to extend our ideas to the case of the harmonically pure mask. 1) Although a sinusoidal mask is mathematically more predictable than a noise mask, it may nonetheless require some practice for the subject to learn the appearance of the sinusoidal mask. Thus a sinusoidal mask may be treated as noise by the visual system until its appearance has become familiar. We suspect that the visual system's natural response to any new masking situation is to make the conservative assumption that some noise (i.e. unpredictability) is present in the mask. Only with practice and pressure to detect with maximal sensitivity will the visual system switch to the more sensitive power law mode of detection. Although we believe this idea to be true, we must acknowledge that it is a conjecture, and that it is not easily tested. 2) A second reason why we may observe Weber's Law with sinusoidal masks is because we find that sinusoidal masks may obey what we term the Extended Birdsall's Theorem. Under appropriate conditions, we find that subjects can use a threshold criterion which amounts to comparing the intensity of the test stimulus with some spatial aspect of the mask stimulus which is proportional to mask contrast. In the strict version of Birdsall's Theorem, this latter was ,the measure of channel output variability and the comparison was done by the critical ratio I/ , which is proportional to I/I_{mean} or Weber's Law. In practice, however, many other geometrical properties of the mask which vary with I would yield the same ultimate ratio. An example of two such criteria can be seen in Figure 5. In the top part of Figure 5 we see sinusoidal mask and test stimuli which have been added to produce a familiar interference pattern. It is readily shown that in the region of destructive interference the width of the grating bars become significantly less than in the region of constructive interference. If the subject adopts as a criterion, let us say, a 15% change in bar width, then this is easily shown to obey Weber's Law with a Weber fraction of about 0.3. This is a plausible criterion; subjects have actually reported to us the they use such a strategy. A second possible criterion is shown in the lower half of Figure 5. In this criterion the subject will set his threshold at the point where the faint dark bar located in the center of the area of destructive interference just disappears. Such a criterion yields a Weber fraction of 1.25. Can a subject actually use such criteria in practice? The results of such an experiment are shown in Figure 6. Here the subject sets his threshold by method of adjustment with a 4 cycle/degree test seen against a 5 cycle/degree mask. In the upper-most curve the subject attempted to use the disappearing-dark-bar criterion. These data obey Weber's Law, and the Weber fraction is approximately as predicted. In the intermediate curve, the subject attempted to use a bar-width criterion. Again it will be seen that Weber's Law is obeyed, and in this case the Weber fraction is about what is seen in other experiments. Finally, in the lower most curve, the subject set his threshold as sensitively as possible, by whatever means he could These thresholds are still lower, and they obey a power law. Thus we see that even with harmonically pure mask stimuli and method of adjustment psychophysics, there are different criteria which a subject may use in setting thresholds, and that these criteria may produce either Weber's Law or power law behavior.

We may draw the following conclusions from this research

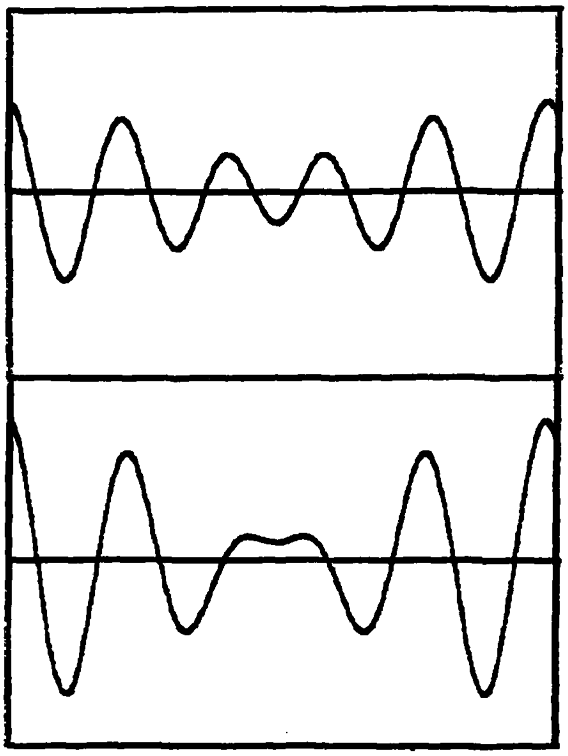

Fig. 5. An explanation of the two adjustment criteria used by the subject in Fig. 6. (See Text.)

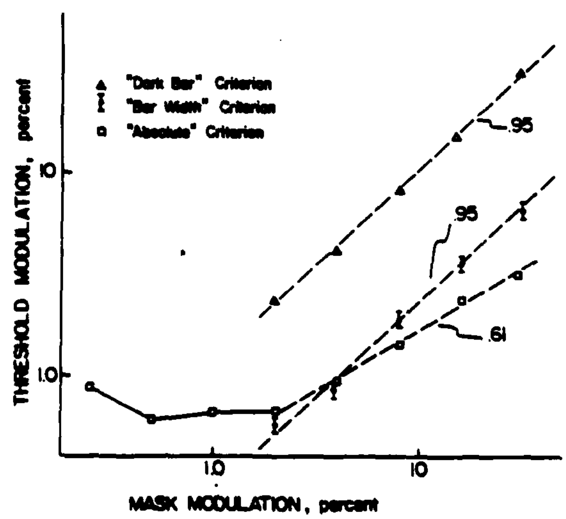

Fig. 6. Masking data (Mask = 5 C/deg, test = 4 C/deg) using the 2 criteria described in Fig 5, as well as an "absolute" criterion. (See text.)

1) Birdsall's Theorem appears to be applicable to the detection of a test stimulus in the presence of visual masking noise. However, what constitutes noise in this regard must be defined subjectively; the same random pattern may function as noise (in the sense of Birdsall's Theorem) when it is unfamiliar, but cease to function as noise when it has become predictable. Moreover, there is a suggestion that even manifestly predictible masks (such as sine waves) may function as noise when they are first seen. A corrollary of this conclusion is that those masking studies which yield Weber's Law behavior are less interesting to a student of the visual system then experiments yielding power law behavior. Birdsall's Theorem shows that the limiting constraints on Weber's Law behavior are probably external to the visual system, while the constraints on power law behavior are presently unknown but are presumed to be internal.

2) Even with harmonically pure masks, there are detection criteria available to the subject which obey either of the two laws. Although it is probably impossible to prove which criterion subjects use in a particular experiment, the existance of such criteria provides a plausible explanation for the appearance of Weber' Law in experiments where noise masking seems inapplicable.

3) At a more general level, it is interesting to note the type of model which we find necessary to account for the results of masking experiments. Spatial frequency detection models in the past have typically utilized simple criteria for detection (excitation in a given channel exceeds some threshold, for example). It is apparent that such simple models are althogether inadaquate to deal with masking phenomena. We have had to postulate processes which approach the realm of cognitive psychology.

4) At a more practical level, it seems possible to use visual masking as a technique for probing the operation of the visual system, particularly its nonlnearties. However, if we are to obtain information about the visual system (and not about external statistical constraints) then the subject must be pushed to detect as sensitively as possible, so the Weber's Law does not apply. Appropriate ways to do this include the use of forced choice physics, with feedback and lots of practice, and expecially the use of three-alternative forced choice which we find to be advantageous with naive subjects.

cc. Sustained and Transient Processing

Results from several areas of research suggest a dissociation in the processing of spatial change (form) and temporal change (motion and flicker). Proposals for such a dissociation originated in the neurophysiological literature, where retinal ganglion cells were classified in accordance with their spatial summation properties (Enroth-Cugell & Robson, 1966) and response characteristics (Cleland, Dubin and Levick, 1971). Shortly thereafter, two-mechanism hypotheses were used in the psychophysical literature to explain various phenomena including different contrast thresholds for the detection of pattern and flicker (Keesey, 1972; Kulikowski & Tolhurst, 1973). More recently, two mechanism hypotheses have been used to account for differences in response to stimuli widely separated in spatial and/or temporal frequency. Much of this work has been reviewed by Legge (1978).

Experiments finding differences between such widely separated stimuli do not require separate spatial and temporal mechanisms for an explanation. Such data could result from a continuum of channels having continuously varying spatial and temporal tunings and no functional segregation related to the perception of spatial and temporal change. This criticism does not apply to two-threshold psychophysical procedures, where testing occurs over a common range of stimulus parameters.

Kulikowski and Tolhurst (1973) used the method-of-adjustment to measure "flicker detection" thresholds and "pattern recognition" thresholds for 0.8 and 12 c/d sinusoidal gratings flickering sinusoidally at temporal frequencies from 0 to 15 Hz. The sensitivity profiles produced by the two thresholds were highly distinctive. Pattern thresholds were low-pass, while flicker thresholds were bandpass. The thresholds were typically separated by a factor of 2 or 3. Harris (1981), using the same procedure with drifting gratings, found a constant ratio between pattern and motion thresholds of gratings drifting at the same velocity. This ratio was precisely proportional to drift velocity. Since velocity may be expressed as ratio of the temporal and spatial frequencies, he proposed velocity to be coded by the ratio of activation of separate temporal-frequency and spatial-frequency sensitive mechanisms. Similar results were obtained using counterphase-flickered gratings for velocity-equivalent conditions.

Other recent experiments have obtained results which conflict with those from the two-threshold procedures. The central issue concerns whether the two criteria actually tap separate spatial and temporal mechanisms or merely reflect different signal strenths within a single perceptual mechanism. Derrington and Henning (1981) and Virsu, Rovamo, Laurinen and Nasanen (1982) used forced-choice procedures to compare the contrast threshold for detecting a counterphase flickering grating with thresholds for discriminating the orientation of the grating. The two functions co-varied over spatial and temporal frequency. Because orientation discrimination is clearly a spatial task, two mechanism models predict that it should require more contrast at low spatial-frequency, where the spatial mechanism is proposed to be least sensitive. Burbeck (1981) used a criterion-free procedure to measure pattern and flicker thresholds in counterphase-flickered gratings. Her results differed substantially from those of the method-of-adjustment procedures, showing that the contrast threshold surface is essentially the same as the pattern-threshold surface alone. Panish (1982) used a multiple-alternative forced-choice procedure to examine the appearance of near-threshold flickering gratings following adaptation. Conditions designed to tap a temporal mechanism showed that a predominence of flicker ratings could be accounted for by an elevated flicker false-alarm rate. In each of the above cases, the use of a less bias-prone technique has significantly reduced those experimental differences which are cited as evidence for a two-mechanism visual system.

There has been a tendency for experimenters to describe the two threshold procedures as different tasks, such as flicker 'detection' and pattern 'recognition.' It may be that these task differences are merely subtle criterion changes. In view of the differences between results from criterion-free/bias-correcting procedures and from method-of-adjustment experiments, it seemed prudent to perform a method-of-adjustment experiment using more carefully described criteria. In addition, a full two-criterion spatio-temporal contrast sensitivity surface has yet to be published and studied in detail. In fact, despite the common assumption that motion and counterphase flicker are largely equivalent, we decided to measure two such surfaces, one for moving and one for flickering gratings.

Experiment 1

The first experiment was designed to study the form and flicker threshold functions over a wide range of stimulus values. Counterphase flickered gratings were chosen to minimize involuntary eye-tracking and afterimages.

Thresholds are plotted in Figure 1. Several striking differences are immediately apparent between these data and those of Kulikowski and Tolhurst (1973). Perhaps the most important is that the shapes of the temporal response functions are generally very similar for the pattern and flicker thresholds, even at very high spatial frequencies. At 8 c/d and 4 c/d both functions are low-pass. At 1 and .5 c/d the curves show more bandpass character. The curves covary consistently for subject SP, while at .5 c/d the threshold functions are distinctly different for subject PC. It is also notable that our data show less than a factor of two separation of the form and flicker curves.

Crossover points between the two functions migrated to

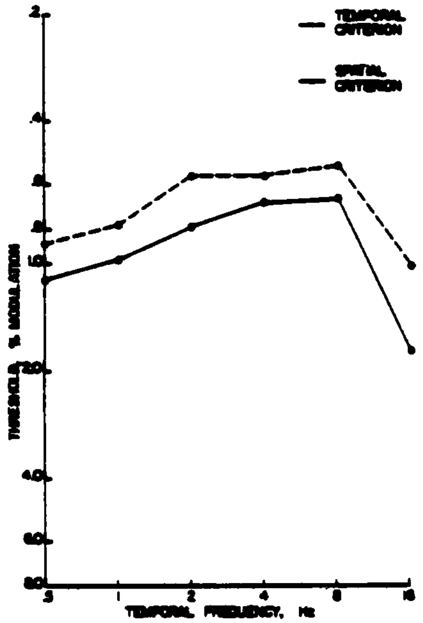

Figure 1-A. Two-criterion contrast thresholds obtained using separate temporal and spatial criteria. Gratings were counterphase-flickered, with a spatial frequency of .5 c/d. Subject SP.

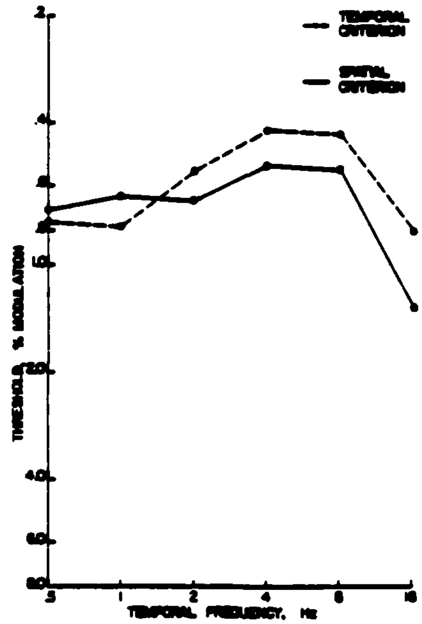

Figure 1-B. As 1-A, except spatial frequency was 1 c/d.

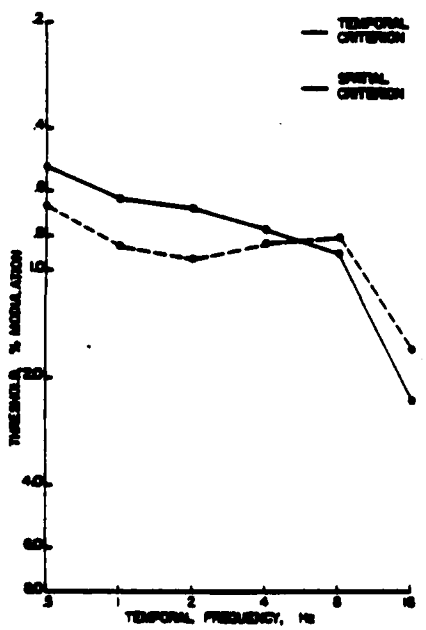

Figure 1-C As 1-A, except spatial frequency was - c/d.

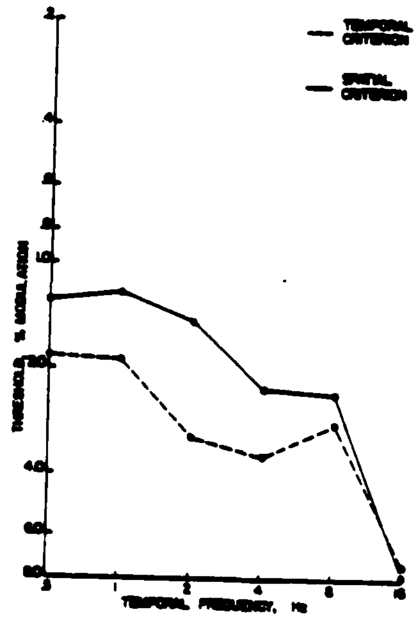

Figure 1-D. As 1-A, except spatial frequency was 8 c/d

higher temporal frequency as spatial frequency increased. The ratios of the two thresholds were not constant for velocity-equivalent conditions, nor did they increase linearly or even monotonically with temporal frequency. Although these data show reliably separable spatial and flicker thresholds, the generally small separation and similar form resemble criterion shifts rather than separate mechanisms. Ratios of temporal and spatial frequency do not appear to offer a basis for the coding of velocity.

Experiment 2

The stimulus and type of temporal modulation used in the two-threshold literature have varied. Keesey (1972) used a counter phase flickering line, Kulikowski & Tolhurst (1973) sinusoidally flickering gratings, and Heeley (1981) drifting and counter phase flickering gratings. In order to test the possibility that the type of temporal modulation could affect the threshold functions, the experiment was repeated using moving, rather than counterphase, gratings.

Notable differences exist between the temporal response functions for moving gratings and those generated by counterphase flickering gratings. Results are plotted in Figures 5-8. These data more closely resemble those obtained by Kulikowski and Tolhurst (1973) than do the counterphase data.

This resemblence is due to the relative invariance of the forms of the temporal response functions across spatial frequency. Flicker thresholds show a clear bandpass function over temporal frequency. The pattern threshold functions are predominantly low pass with some bandpass character, particularly at low spatial frequency. Crossover points of the two functions rise with spatial frequency. Ratios of the temporal to spatial thresholds varied monotonically with spatial frequency and inversely with temporal frequency. The slopes of the ratios across velocity range between .31 and .47.

It is clear that subjects can consistently set separate thresholds for the perception of temporal and spatial variation, and that there is a general correspondence between data from different subjects. However, the divergence in the data from the two temporal conditions presents several problems: 1) can either condition be regarded as more fundamental, explaining the other, 2) why do these results differ from those of earlier experiments, 3) can these results support a two mechanism model of spatial and temporal vision?

No obvious transform relates the results obtained from the drifting and counterphase conditions. They may represent different visual processes. However, we speculate that the differences between the two data sets result from two factors: the use of different visual strategies for optimizing the desired percepts, and the difficulty of avoiding tracking with drifting

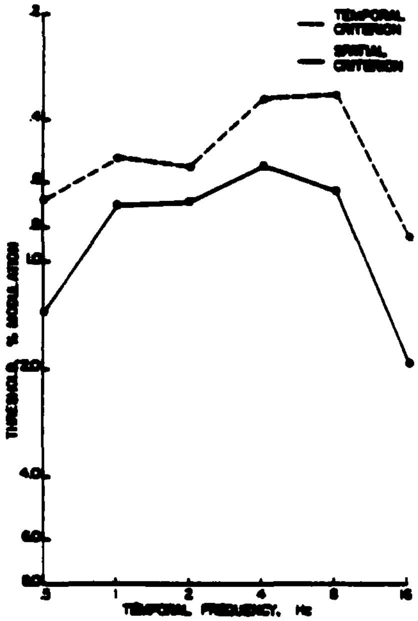

Figure 5-A. Two-criterion contrast thresholds obtained using moving gratings. Spatial frequency was .5 c/d. Subject SP.

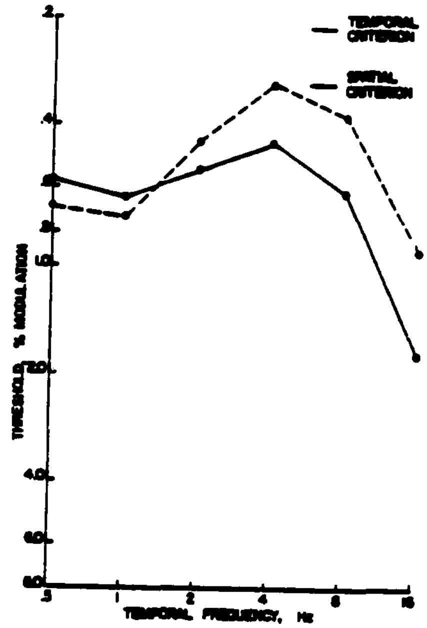

Figure 5-B. As 5-A, except spatial frequency was 1 c/d

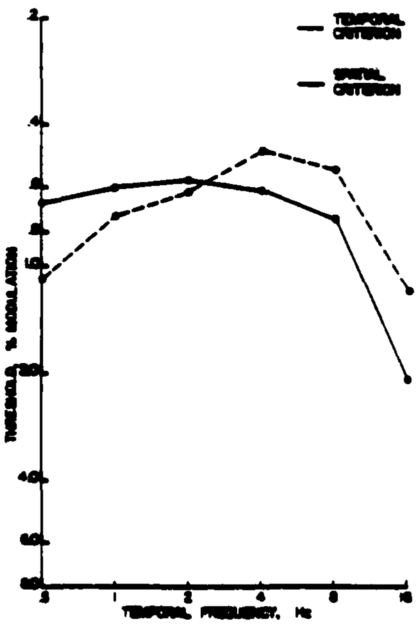

Figure 5-C. As 5-A, except spatial frequency was [illegible] c/d.

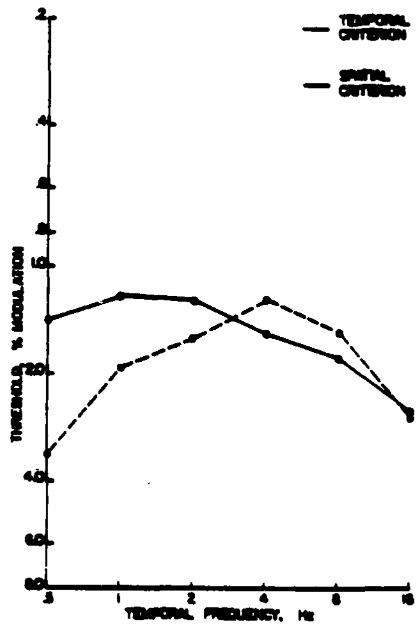

Figure 5-D. As 5-A, except spatial frequency was 8 c/d.

gratings. Kelly (1982) has demonstrated that thresholds for unstabilized counterphased gratings closely approximate results obtained with stabilization. Experiments which use drifting gratings are more prone to artifacts resulting from involuntary eye-movements (tracking). Since tracking effects the temporal parameters of the retinal image, it is an important factor in the measurement of thresholds, particularly at high spatial frequencies. (Arend, 1976). The net effect of tracking is to reduce the effective temporal frequency of the stimulus. This would result in less high temporal frequency fall off when using drifting gratings, particularly at high spatial frequencies. This may be observed clearly in the pattern threshold functions less so in the motion threshold functions. This difference may result from use of different patterns of eye-movement in the two tasks, such that adopting a pattern criterion results in more tracking than does a motion criterion. If such eye-movements cause the differences observed between the counterphase and drift conditions, the differences should be abolished by retinal stabilization of the stimuli. We will be performing such experiments in the near future.

The results from the counterphase condition in this study differed substantially from those of Kulikowski and Tolhurst (1973) and Harris (1981). However, our data are very similar to those from the two method-of-adjustment conditions of Burbeck (1981). The reason for these differences is unclear. However, the two-mechanism hypothesis was supported by the independence of the pattern and flicker threshold. Pattern and flicker thresholds obtained in the present study covary considerably, although they are clearly separable. This separability could result solely from criterion shifts, without involvement of separate mechanism. These data together with the results of Derrington and Henning (1981), Burbeck (1981) and Panish (1982) suggest the need for a careful reconsideration of the evidence for separate pattern and flicker mechanisms.

dd. Channel Tuning for Motion

Our major study on motion tuned channels was a substantial parametric study of the motion aftereffect. In this study we considered the effect of adapting to a moving sinusoidal grating upon the perceived of velocity of a second moving grating. There are at least four significant parameters in this study: these being the velocity and spatial frequency of the mask and test gratings. We studied the two velocities in depth, using the same spatial frequency for both gratings, as well as taking some partial data at other spatial frequencies. Velocity matching is a tedious psychophysical task, and these studies involved several hundred hours of subject time. The results are extremely complicated, and cannot as yet be said to have simplified our understanding of the motion system. One result is absolutely clear: the MAE is in no sense a simple velocity which is algebratically added to the test velocity. The MAE is strongly effected by the relative difference in adapt and test velocities. Our data suggest a sort of velocity shift effect (analogous to the spatial-frequency shift) in which the deviation of the test velocity can be in either direction, depending upon relative velocities of adapt and test. Super-imposed upon this, however, are other effects which we have not yet characterized, making the whole picture extremely complex. As we were about to prepare a more detailed write up of this experiment, we found a paper by Thompson (1982) which performs essentially similar experiments. The major difference in our results is that Thompson failed to find any indication that the motion after effect could change perceived velocity in both directions. In every case his motion after effect was in the direction opposite to the adapt motion, as has been traditionally reported. We are now attempting to replicate our findings with more subjects (including naive ones).

ee. Image Enhancement

We have studied the visual system's inherently poor sensitivity to low frequency targets, and some potential ways to circumvent this. There is already a standard technique for enhancing the visibility of low spatial frequencies: that is the mapping of grey scale values onto a full spectrum of pseudo-colors. This technique is undeniably successful, though some training is required to interpret the resulting images. A more serious difficulty with this technique is that it can only be used with monochrome images, which are then presented in full color. It cannot be used with natural color images, nor with monochrome images presented in monochrome. It is our hope to develop enhancement techniques which avoid these shortcomings. The basis for much of our work is shown in figure 7. This is a picture of the two dimensional spatio-temporal contrast sensitivity surface, which was taken from a paper by Kelly (1978). It will be seen that when both spatial and temporal frequency are low, then sensitivity is also very low. However even if spatial frequency is held low detectability can be made

Visual contrast sensitivity

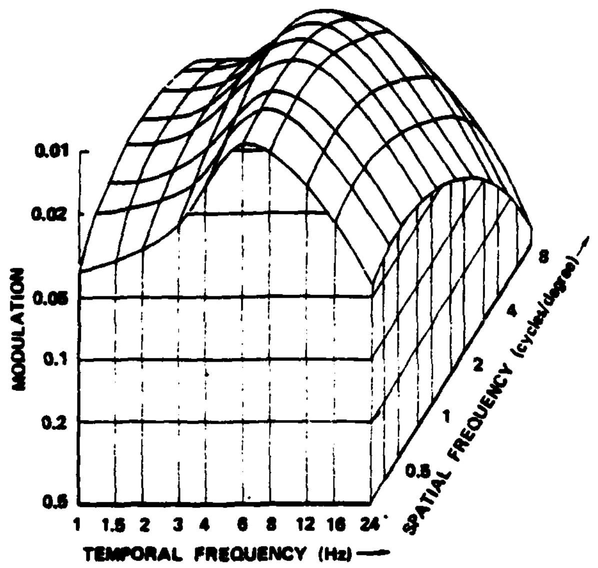

Fig. 7. Spatio-temporal threshold surface, showing the effects of combined spatial and temporal sine-wave modulation on the contrast sensitivity.

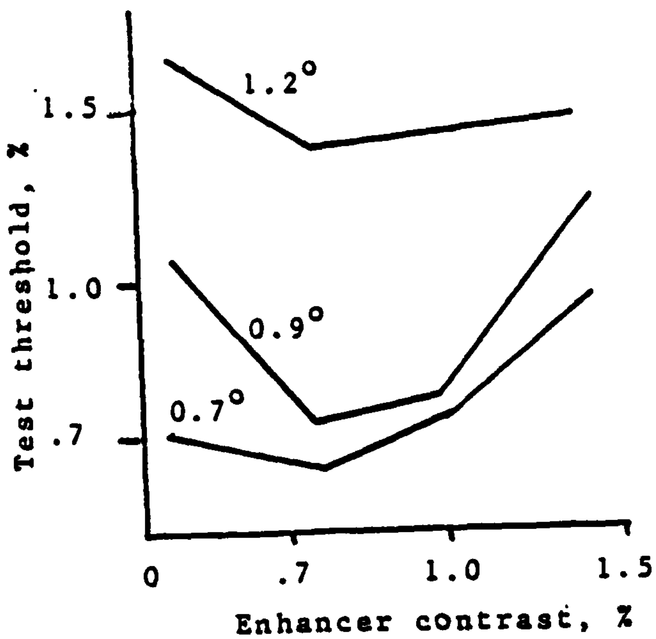

Fig. 8. Threshold for a Gausslan bar as a function of enhancer contrast for 3 bar SD's. Enhancer is 0.5 c/deg.

to increase enormously by raising the temporal frequency. Thus it should be possible to enhance the visibility of a low frequency image by adding to it some temporal variability.

A demonstration of this effect is shown in the first part of the accompanying video tape. Here we see the faint image of a human face, which is almost entirely lost in white noise. We now use the very simple technique of rapidly reversing the contrast of this image, providing some counter-phase temporal modulation, and you will observe that the image appears much more detectable. This is a strong effect; it has been studied by a number of investigators and a sensitivity improvement of as much as a log unit may occur with very low spatial frequencies.

We may conceptually divide image enhancement schemes into two classes, the linear enhancers and the non-linear enhancers. Linear enhancers are based directly upon the shape of the two dimensional contrast sensitivity surface shown in figure 7, and upon the assumption that contrast detection is a reasonably linear process. The existence of non-linear enhancers has not yet been demonstrated, but the hope is that procedures exist which will produce a supra-ordinate enhancement effect. That is, one which is either greater than that predicted by figure 7 or an enhancement effect which is qualitatively superior (an example might be an enhancer which works better with natural stimuli than with sine waves.)

Another example of a linear enhancer is shown in the next demonstration. Temporal variation can be added to an image either by flickering the image, as in the previous demonstration, or by making the image move. This demonstration shows a gaussian bar which is subthreshold when held stationary, but which is very easily seen when moved across the screen. Some data on this effect are shown in figure 10. Unlike the demonstration, in which the display was noise-free, the data in figure 10 are for an image which is noise limited. It is clear that this enhancement technique works well in the presence of noise.

It is appropriate here to say a few words about the two types of noise which are used in figure 10; these are static and dynamic noise. Static noise does not change with time, and is simply a part of the image. Such noise is typical of aerial photographs and other one-time images. Dynamic noise changes randomly from frame to frame. This type of noise is typical of real-time displays.

We have studied two different types of potential non-linear enhancers. The first is based on the well known observation that the presence of a subthreshold signal may enhance the detection of a second test signal superimposed upon the first. This has been studied in some detail with sinusoidal enhancers, using either other sinusoids or luminous lines as targets. With such stimuli, target detectability may be more than doubled. A problem with this technique is that the test stimulus may be enhanced or "de-enhanced" depending upon whether it falls on a

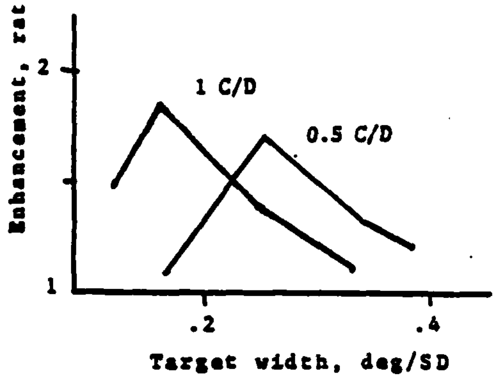

Fig. 9. Enhancement as a function of bar size for a 0.75% contrast, 0.5 c/deg enhancer.

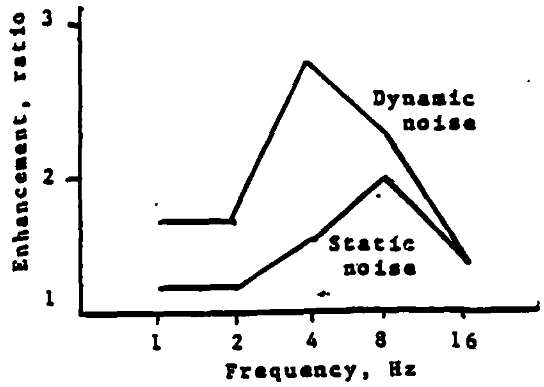

Fig. 10. Enhancement proceded by moving the whole field back and forth at the indicated frequency. Target is Gaussian bar, 0.8 /SD.

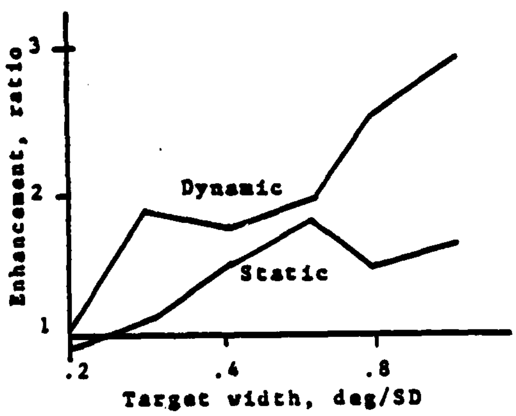

Fig. 11. Enhancement produced by Pixel flicker (1Hz) for static and dynamic noise. Target is Gaussian bar of different widths.

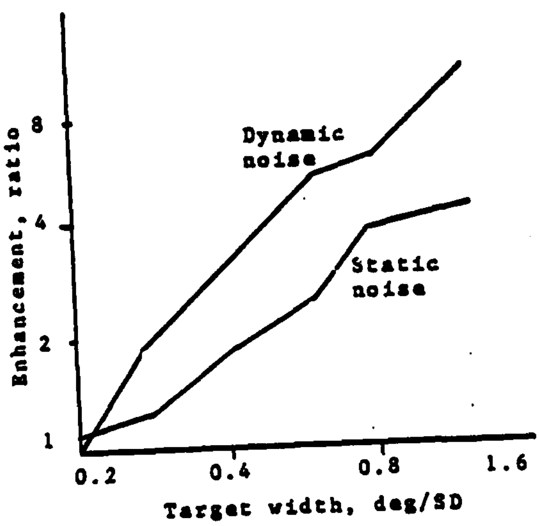

Figure 12, Pixel Flicker enhancement. Same as Fig. 11, except there is a dark bar adjacent to the target.

maximum or a minimum of the enhancing grating. We have tried to deal with this by having the gratings drift slowly across the display so that all portions of the display are enhanced at some time. Some results are shown in figure 8 for a gaussian bar with a sinusoidal enhancer. For an appropriate enhancer contrast, it is clear that significant enhancement occurs. Figure 9 shows the effect of the relative sizes of enhancer and target. Sinusoidal enhancers are apparently most effective with targets whose dimensions are comparable to the half wave-length of the sinusoid.

A final type of non-linear enhancer which we have studied in some detail is Pixel Flicker (Genter and Weisstein 1980). Although Weisstein claimed that Pixel Flicker was a true supra-ordinate enhancer, having a greater effect upon natural stimuli, our initial studies were necessarily done with gaussian bar targets. Although these are hardly natural stimuli, the results are nonetheless encouraging.

Briefly, Pixel Flicker works as follows. Consider a two-dimensional raster image, digitized on a grey scale ranging from 0 to 255, which is being displayed at 30 frames/second. During the first frame, the image is displayed unmodified. During the second frame, a constant (1, for example) is added to each brightness value. Most of the pixels simply get brighter, but those which had a brightness of 255 in the original image now become 0 (modulo 256 arithmetic), producing a paradoxical brightness reversal. In each successive frame, the constant is again added to each brightness value, so the the original image is restored after 256 frames, with every brightness value having cycled through the reversal and back to its original value. The resultant, time-varying image is -- to say the least -- bizarre, but what is totally unexpected is the effect of this procedure in the presence of static noise. Genter and Weisstein added random numbers, with a range of 0 to 200, to every pixel in their image, making the original image essentially invisible. When the Pixel Flicker technique is applied, however, the image becomes dramatically visible.

We have systematically studied Pixel Flicker in one dimension. The enhancement ability of Pixel Flicker with noise-free images is theoretically infinite; as a result we have concentrated on the static and dynamic noise conditions. This is a complicated study, and our present results can only be considered indications. In figure 11 we see the amount of enhancement as a function of the size of the target (which is a Gaussian bar). It is clear that enhancement increases as the bar becomes larger (lower spatial frequencies) and that Pixel Flicker is more effective with dynamic, rather than static, noise. Although the actual amount of enhancement varies considerably with conditions (the largest effect we have seen is about 10x enhancement), these two general conclusions seem to hold. It became clear to us quite early in these experiments that if our larger Gaussian bars were placed too close to the edge of the display, then they were very effectively masked by the dark surrounding area. This is exactly as predicted by spatial masking. We tested this more carefully by placing a dark bar (a cardboard mask) in the middle of the screen. The test bars appeared on either side of this bar (we were using forced-choice psychophysics, with the targets displayed randomly on the left and right sides of the CRT). The results are shown in figure 12. It is clear that Pixel Flicker is especially effective in enhancing masked images. In figure 13 we plot the amount of enhancement as a function of the Pixel Flicker rate. This is the rate at which the brightness of the display cycles once through the reversal point; it is directly proportional to the constant which is added to the brightness after each frame. We see that enhancement is maximum at a flicker rate of about 1-2 Hz. Weisstein reports a similar finding.

About a month before the 1982 AFOSR meeting, our image processor became functional and we were able to perform some qualitative studies of Pixel Flicker with real images. Our first demonstration is a nearly exact replication of the demonstration presented by Genter and Weisstein at the OSA meeting in 1980. There is one unusual aspect of this demonstration and that is the peculiar way in which noise was added to the image. Genter and Weisstein started with an image having a brightness range of 0-255, and added to each pixel uniformly distributed random numbers with a range of 0-200. When the sum exceeded 255, the result was taken modulo 256. Thus many of the brighter points in the original image are mapped into dark points in the noisy image. This modulo-noise is not found in functioning visual displays; normally the display gain would be adjusted to minimize the number of points exceeding the dynamic range of the display, and those points which did exceed would be clamped at the maximum value. In fact, we shall see that modulo-noise degrades the image far more than a comparable amount of ordinary noise, and that the alleged supra-ordinate effect of Pixel Flicker does nothing more than compensate for some of this loss.

In the original demonstration of pixel flicker, the ratio of signal-to-noise amplitudes was approximate 1:1. The next demonstration shows the same image with a 1:1 signal-to-noise ratio, but with the signal and noise amplitudes reduced so that no pixel can exceed 255. Despite the equal nominal signal-to-noise ratio, it is clear that the second demonstration is enormously more visible than the first. The next demonstration shows the same image using natural noise, with a signal-to-noise ratio of about 1:8. The visibility of this image seems comparable to that of the image with modulo-noise at a signal-to-noise ratio of 1:1. Why is it that modulo-noise is so much more effective in degrading an image? Consider Weisstein's original modulo noise image. A pixel with a brightness value of 0 in the noise free image will have a value uniformly distributed in the range of 0-200 in the noisy image. Since the mean of the added noise is 100, then mean brightness of the noisy pixel will also be 100. In general, for noise-free pixels in the range of 0-55 their mean transformed brightness values will be 100-155 respectively. However, if the noise-free pixel is 56 or greater, then the possibility exists that the noise-added pixel will

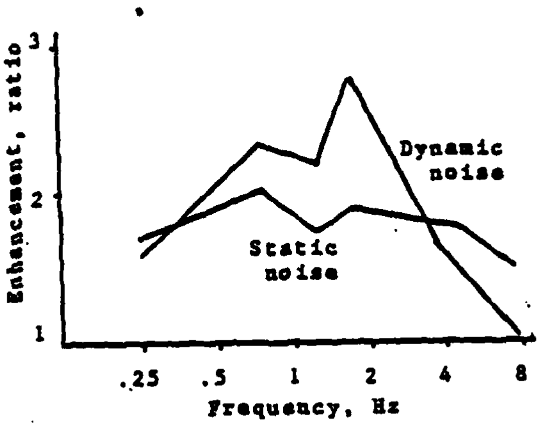

Fig. 13. Pixel flicker enhancement of 0.6 deg/SD Gaussian bar for different cycling rates.

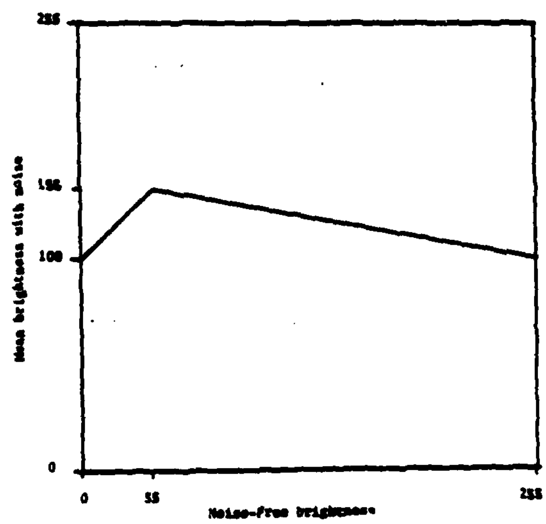

Fig. 14. Mean brightness of image points before and after the addition of modulo-noise (noise range 0-200, uniformly distributed).

exceed 255 and the actual sum, modulo 256, will be close to 0. The addition of these small values actually causes the mean brightness of the noise-added pixel to decrease as the noise-free value increases. The net result of all of this is shown in figure 14 which plots the brightness of a noise-free pixel against its mean brightness after the addition of noise. It will be seen from the leftmost branch of this figure that those parts of the image which have brightnesses in the range of 0-55 are maped linearly into the range of 100-155. An image which contains only these brightness values would be only slightly degraded; it receives an additive background plus some additive noise. The situation is much worse for brightnesses in the range of 56-255. Unlike the first branch of the curve, with a slop of 1, this branch of the curve has a slope of approximately minus 0.25. Thus contrast values in this range are not only reduced four-fold, but are actually made negative. A close examination of Weisstein's original demonstration confirms this; it is now seen that most of the image is indeed at reduced negative contrast.

How does the Pixel Flicker algorithm render this image more visible? It is readily shown that as the Pixel Flicker procedure cycles through the entire range of brightnesses, the entire curve in figure 10 translates laterally. Thus the smaller branch of the curve which has unity contrast gain passes cyclically through every range of brightness. Thus each region of brightness experiences a fraction of a Pixel Flicker cycle in which it is reproduced with minimal distortion. Close examination of the Pixel Flicker demonstration shows that this agrees with subjective impression; details pop out of the noise for a brief portion of each cycle and the disappear. The point can be proven by running Pixel Flicker very slowly through its cycle, and stopping it at appropriately selected points. For example when stopped about 2/3 of the way through a cycle, some of the brighter portions of the image (notably the face) become highly visible, although darker portions are lost. The fact the Pixel Flicker works even when the algorithm is stopped demonstates that this is not a temporal phenomenon at all, as originally claimed by Genter and Weisstein.

In a final demonstration, we apply Pixel Flicker on an image with natural noise rather than modulo noise. As noted above, it is necessary to use a great deal more natural noise in order to obscure the image to a comparable degree. When this is done, the results are frankly disappointing. The dramatic effects of the original Pixel Flicker demonstration are entirely absent. In fact it is not apparent that there is any enhancement whatsoever. certainly the demonstration is less effective than any of the preceeding ones.

We may draw the following conclusions about Pixel Flicker as an enhancer. 1. Our studies with gaussian bars show Pixel Flicker is useful in extracting low frequency images from noise. It is unclear whether it is superior to other, simpler techniques. 2. Contrary to the claims of Genter and Weisstein, there is no evidence that Pixel Flicker possesses supra-ordinate, non-linear enhancing properties when used with natural stimuli. The demonstrations which would seem to show such effects have been explained in terms of the manner in which noise was added to the stimuli. These effects disappear when more realistic noise is used. Finally, these effects are not dependent upon the temporal properties of pixel flicker; they work equally well with a static transformation.

BIBLIOGRAPHY

Burbeck, C. Criterion-free pattern and flicker thresholds. Journal of the Optical Society of America, 19 81, 71, 1343-1350.

Derrington, A and Henning, G. Pattern discrimination with flickering stimuli. Vision Research, 1981, 21, 597-602.

Enroth-Cugell, C. S. Robson, J. The contrast sensitivity of retinal ganglion cells of the cat. Journal of Physiology, 1966, 187, 517-552.

Green, M. Psychophysical relationships among mechanisms sensitive to pattern, motion, and flicker. Vision Research, 1981, 21, 971-983.

Kulikowski, J. & Tolhurst, D. Psychophysical evidence for sustained and transient detectors in human vision. Journal of Physiology, 1973, 232, 149-162.

Legge, G. Sustained and transient mechanisms in human vision: Temporal and spatial properties. Vision Research, 1978, 18, 69-81.

Panish, S. Time-course of spatio-temporal adaptation; with respect to the existence of separate spatial and temporal visual mechanisms. Unpublished Doctoral Dissertation, University of New Hampshire, 1982.

Smith, R. Studies of temporal frequency adaptation in visual contrast sensitivity. Journal of Physiology, 1971, 216, 531-552.

Swift D. & Smith, R. An action spectrum for spatial-frequency adaptation. Vision Research, 1982, 22, 235-246.

Tolhurst, D. Sharpe, C., & Hart, G. The analysis of the drift rate of moving sinusoidal gratings. Vision Research, 1973, 13, 2545-2555.

Tulunay Keesey, U. Flicker and Pattern Detections: A Comparisen of Thresholds. Journal of the Optical Society of America, 1972, 62(3), 446-448.

Virsu, V., Rovamo, J., Laurinen, P. & Nasanen, R. Temporal contrast sensitivity and cortical magnification. Vision Research, 1982, 22, 1211-1217.

Watson A. and Nachmias, J. Patterns of temporal interaction in the detection of gratings. Vision Research, 1977, 17, 893-902.

Wilson H, and Bergen, J. A four-mechanism model for threshold

d. Publications

Swift, Dan J. and Smith, Robert A. The Action Spectrum for Spatial Frequency Adaptation. Vision Research, Feb 1982.

Swift, Dan J. and Smith, Robert A. Spatial Frequency Masking and Weber's Law. Accepted, Vision Research.

Smith, Robert A. Size Discrimination with Low Spatial Frequencies. Accepted, Perception.

Swift, Dan J., Panish, Steve C.and Smith, Robert A. A Quantitiative Study of the Motion After-Effect. In preparation.

Smith, Robert A, and Swift, Dan J. The Application of Birdsall's Theorem to Spatial Frequency Masking. In preparation.

Panish, S., Smith, R., Swift, D. Pattern and Flicker Threshold in counterphase flickering and moving gratings; separate mechanisms or criteria shifts? In preparation.

Panish, S. Smith, R. and Swift, D. Time-course of spatio-temporal adaptation with respect to the existence if separate spatial and temporal visual mechanisms. In preparation.

e. Professional Personnel

Principal Investigator -- Robert A. Smith, PhD
Research Scientist -- Dan J. Swift, PhD (Summer '82)
Reserach Scientiest -- Larry Arend, PhD (9/1/81 through 8/1/82)

f. Interactions

Swift, Dan J. and Smith, Robert A. Spatial Frequency Masking and Birdsall's Thorem. Present to ARVO, May 1982

Smith, Robert A. Size Discrimination at Low Spatial Frequencies. Presenteds to ARVO May 1982.

Smith, Robert A. Spatial Frequency Masking and Visibility. Prewsented to AFOSR May 1982.